

FIG. 1




Fig. 2




Fig. 3




Fig. 4A

LJPU Software Diagram

350

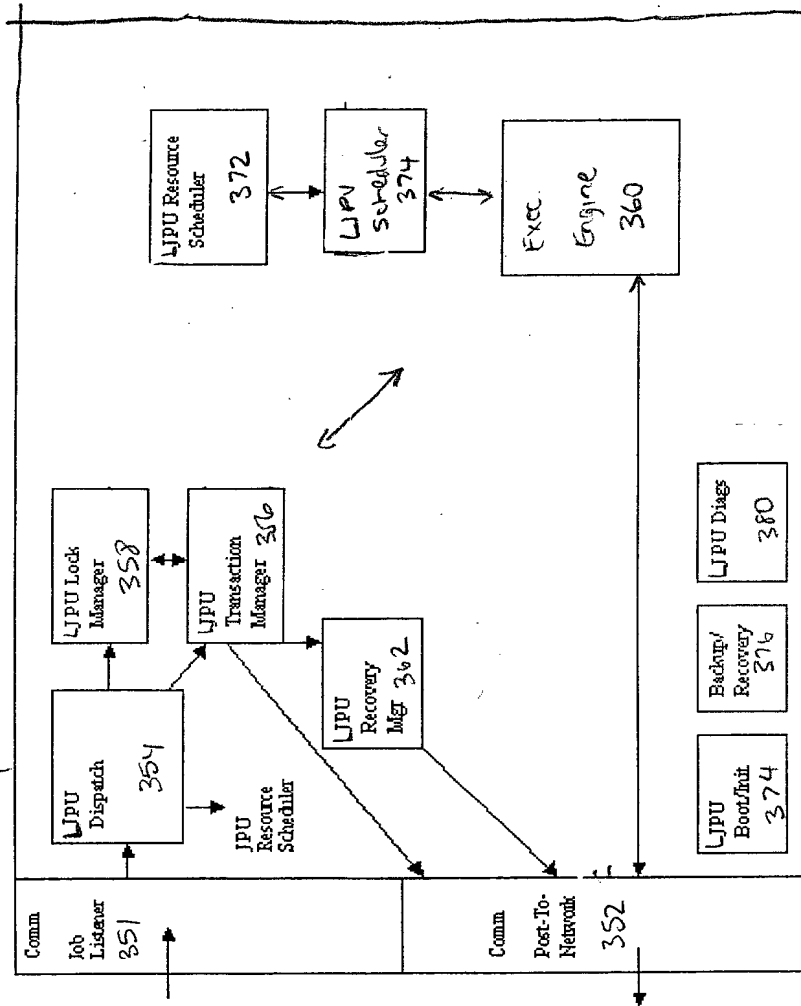


Fig. 4B

28




Fig. 5

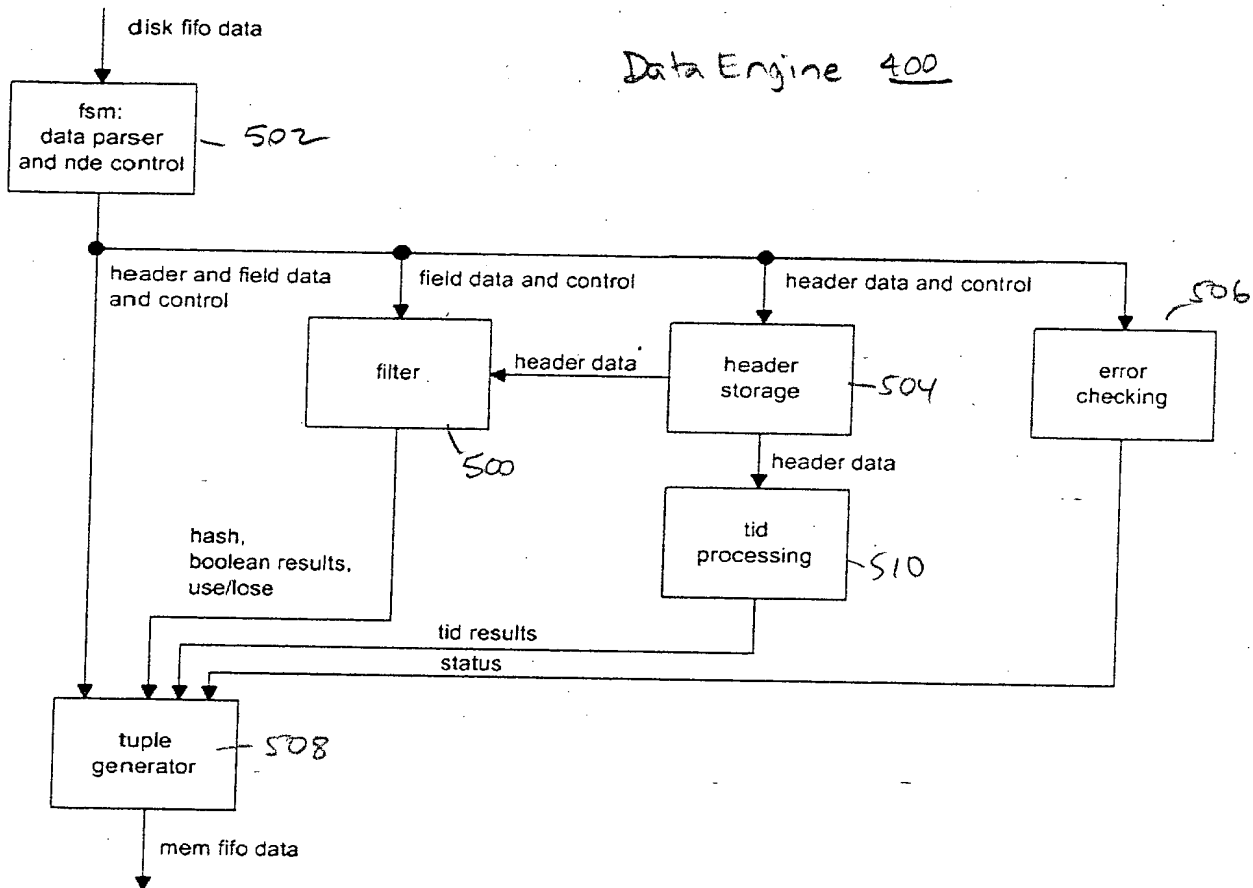


Fig. 6


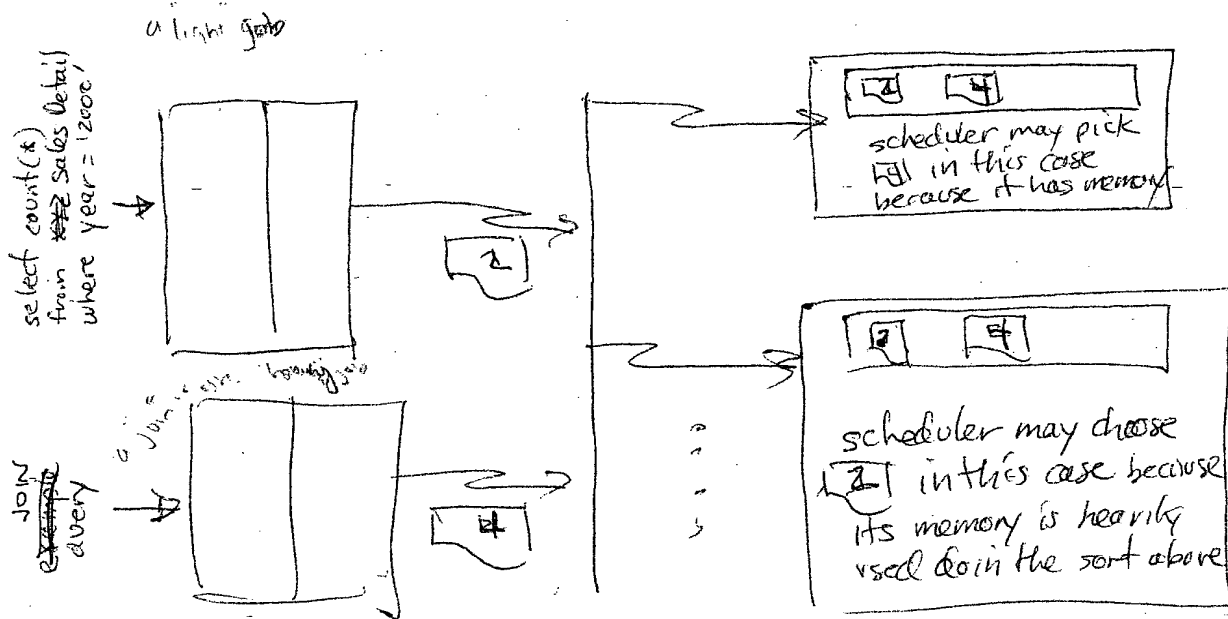
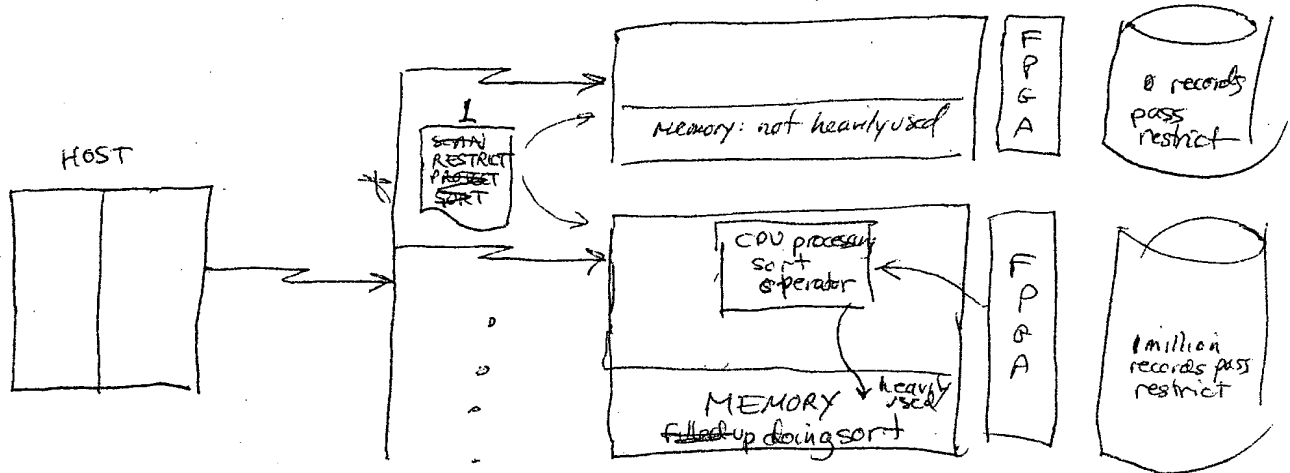


Fig. 7





1 → Job 1: the predicate "year = 2000" can be processed within the FPGA. Doing the count does not require continuing presence of each record. So it does not use much memory (Each record that passes the FPGA filter can be discarded once it is counted)

4 → Job 4: (or 5 or 6) is potentially memory intensive as it is joining 2 streams, one of which must be memory resident

Fig. 8